



US009317964B2

(12) **United States Patent**
Akenine-Moller et al.

(10) **Patent No.:** **US 9,317,964 B2**
(45) **Date of Patent:** ***Apr. 19, 2016**

(54) **DEPTH OF FIELD RASTERIZATION**

(71) Applicant: **Intel Corporation**, Santa Clara, CA
(US)

(72) Inventors: **Tomas G. Akenine-Moller**, Lund (SE);
Carl J. Munkberg, Malmo (SE); **Jon N.**
Hasselgren, Bunkerflostrand (SE);
Robert M. Toth, Lund (SE)

(73) Assignee: **Intel Corporation**, Santa Clara, CA
(US)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 0 days.

This patent is subject to a terminal dis-
claimer.

(21) Appl. No.: **14/688,107**

(22) Filed: **Apr. 16, 2015**

(65) **Prior Publication Data**

US 2015/0287238 A1 Oct. 8, 2015

Related U.S. Application Data

(63) Continuation of application No. 13/096,165, filed on
Apr. 28, 2011, now Pat. No. 9,058,697.

(51) **Int. Cl.**
G06T 15/40 (2011.01)
G06T 17/20 (2006.01)

(52) **U.S. Cl.**
CPC **G06T 15/40** (2013.01); **G06T 17/20**
(2013.01)

(58) **Field of Classification Search**

None

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,359,629 B1 * 3/2002 Hopcroft G06T 15/40
345/423
6,480,205 B1 * 11/2002 Greene G06T 15/405
345/631
6,573,895 B1 * 6/2003 Carter G06T 15/40
345/423
7,812,837 B2 * 10/2010 Naoi G06T 15/405
345/421

OTHER PUBLICATIONS

PCT International Search Report and Written Opinion issued in
corresponding PCT/US2012/033979 dated Nov. 28, 2012 (8 pages).
Office Action [w/English translation] dated Jul. 11, 2014 by the
Taiwan Intellectual Property Office in Application No. 101114149 (4
pages).

Fatahalian, K., et al., Data-Parallel Rasterization of Micropolygons
with Defocus and Motion Blur, Proceedings of High Performance
Graphics, 2009, 10 pages.

* cited by examiner


Primary Examiner — Said Broome

(74) *Attorney, Agent, or Firm* — Trop, Pruner & Hu, P.C.

(57) **ABSTRACT**

Depth of field may be rasterized by culling half-space regions
on a lens from which a triangle to be rendered is not visible.
Then, inside tests are only performed on the remaining
unculled half-space regions. Separating planes between the
triangle to be rendered and the tile being processed can be
used to define the half-space regions.

30 Claims, 7 Drawing Sheets






FIG. 1








FIG. 3










FIG. 7





FIG. 8



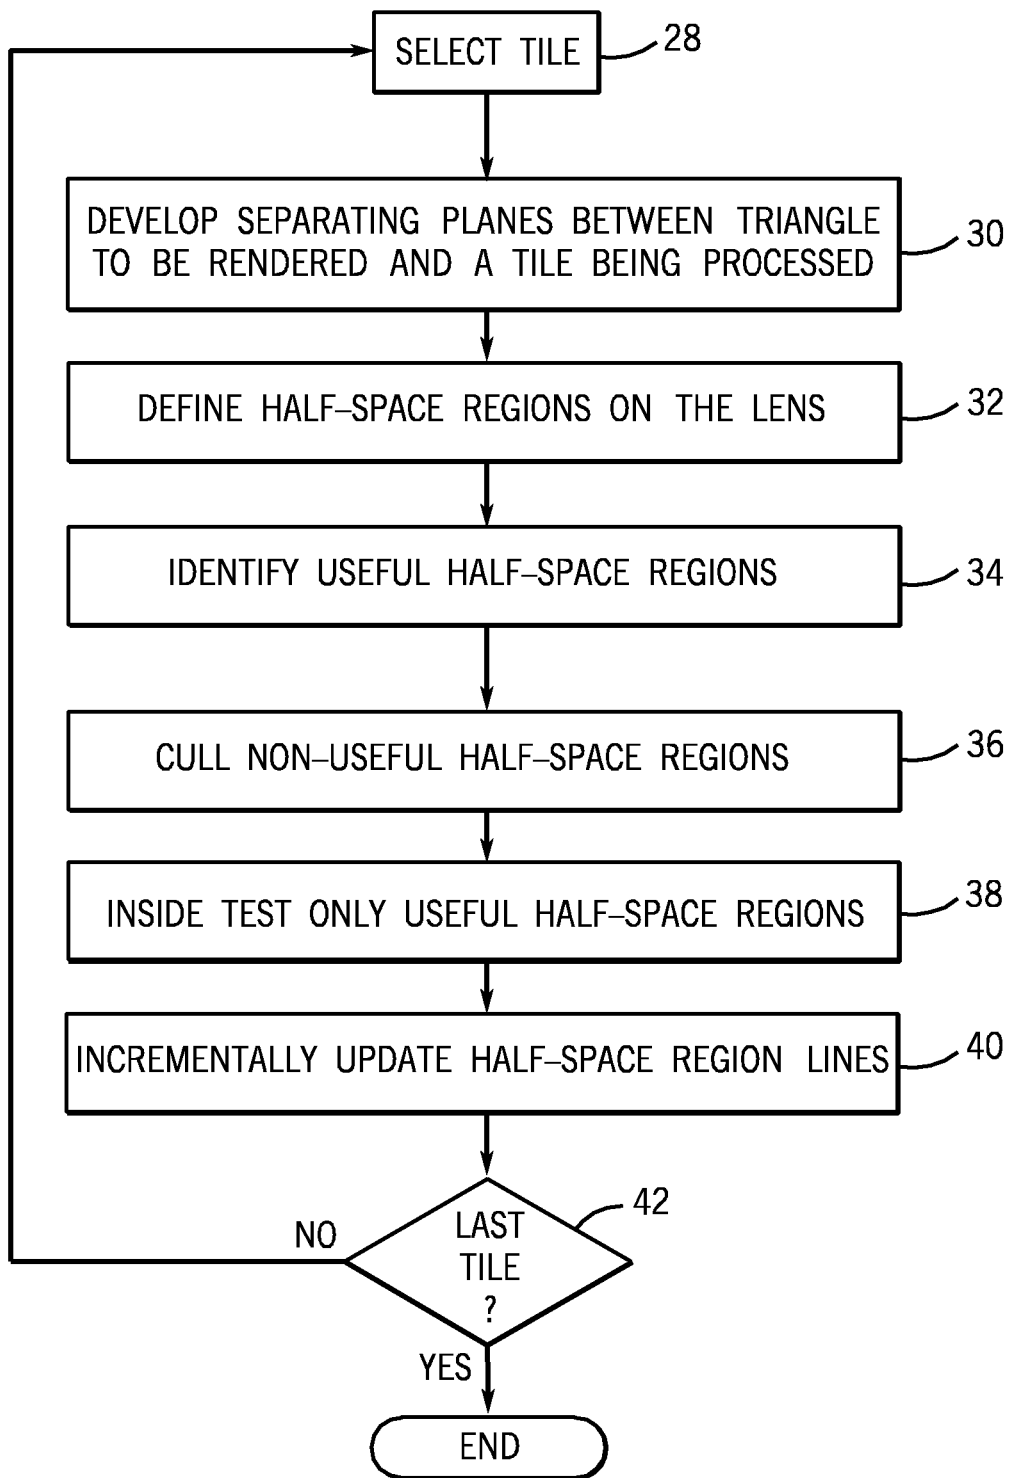


FIG. 10




FIG. 11

1

DEPTH OF FIELD RASTERIZATION

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 13/096,165, filed on Apr. 28, 2011.

BACKGROUND

This relates to computers and, particularly, to graphics processing.

During the last few years, research activity has increased on stochastic rasterization. Efficient rasterization of depth of field (DOF) and motion blur at the same time remains an elusive goal. Rasterization of only depth of field is a substantially simpler problem, but still, specialized algorithms for this are not well explored.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a sketch showing a separating plane in two dimensions, according to one embodiment;

FIG. 2 is a schematic depiction of potentially visible and not visible half-spaces, according to one embodiment;

FIG. 3 is a depiction of how separating planes generate half-space lines, according to one embodiment;

FIG. 4 is a schematic depiction of a triangle projected through tile corners according to one embodiment;

FIG. 5 is a sketch showing visible regions for a triangle seen through a tile in accordance with one embodiment;

FIG. 6 is a depiction of a triangle intersecting a focus plane, according to one embodiment;

FIG. 7 is a sketch showing a triangle edge qr as an intersection tested against the focus plane which generates the intersection point v in accordance with one embodiment;

FIG. 8 is a sketch showing two candidate planes dividing space into four regions in accordance with one embodiment;

FIG. 9 is a sketch of a lens grid with corresponding storage;

FIG. 10 is a flow chart for one embodiment; and

FIG. 11 is a hardware depiction for one embodiment.

DETAILED DESCRIPTION

An efficient algorithm for rasterization of depth of field is based on removing half-space regions of the (u, v) -space on the lens from where the triangle to be rendered definitely cannot be seen through any pixel within a tile. We find the intersection of all such regions exactly, and the resulting region is used to reduce the number of samples in a tile where the full "sample-inside-triangle" test needs to be performed. We present several efficient methods on how to use the resulting region in (u, v) -space to quickly cull samples from being tested.

As usual for depth of field (DOF) rendering, the lens area is parameterized by $(u, v) \in \Omega \subseteq [0, 1] \times [0, 1]$, where Ω is the aperture shape and may, for instance, be square or circular. In general, we have n samples per pixel for stochastic rasterization, where each sample consists of a spatial position, (x_i, y_i) , and a lens position, (u_i, v_i) . A clip-space vertex of a triangle is denoted $p^i = (p_x^i, p_y^i, p_z^i, p_w^i)$, and a triangle is then $p^0 p^1 p^2$. A tile is a rectangular block of pixels in the plane in focus. Sometimes, we use the term "focus plane" instead of "plane in focus," which is the plane with $w=F$ where rendered geometry will appear in perfect focus. An inside test simply computes whether a certain sample, (x_i, y_i, u_i, v_i) , is inside the tri-

2

angle being rasterized. In general, the number of unnecessary inside tests should be minimized.

In order to describe our algorithm, we will use the notion of separating planes, which has been used for cell-based occlusion culling for large occluders. A separating plane between two convex polyhedral objects is a plane formed by an edge from one object and a vertex from the other object, and at the same time, the objects must lie on opposite sides of the plane. An illustration of this is shown in FIG. 1 in two dimensions.

In FIG. 1, the separating lines of a tile and a triangle can be used to find regions on the lens where we can reduce the number of computations needed for stochastic rasterization. The regions 10 and 12 on the lens are detected using the separating planes. Note that in the regions 12, there are no lens positions that can "see" the triangle through any point in the tile. On the left in FIG. 1, a triangle is behind the focus plane. On the right in FIG. 1, a triangle is in front of the focus plane. When a vertex behind the focus plane is used to define the separating plane, the positive half-space contains the tile, and vice versa.

Separating planes, derived from a tile and a triangle, may be used to remove half-space regions on the lens from further processing. We find regions on the lens which cannot "see" the triangle through any point in the tile, and hence all samples in the tile with lens coordinates, (u_i, v_i) , in such regions do not need any further inside-testing. Intuitively, this can be understood from FIG. 1, where the separating lines of the tile and the triangle are used to cull regions 12 on the lens.

In three dimensions, there are two different types of separating planes that can be used to cull half-space regions on the lens. These are illustrated in FIG. 2. The first set of separating planes are generated by a tile edge and a triangle vertex. Let us denote these planes by π_i , where the positive half-space of the plane consists of all points, p , such that $\pi_i(p) \geq 0$. Now, consider the leftmost (vertical) edge of the tile. In the example to the left in FIG. 2, the tile's left side creates a separating plane with the rightmost triangle vertex. This separating plane cuts the lens area into two half-spaces. We call the dividing line a half-space line, $h_i(u, v) = 0$.

In FIG. 2, on the left, in this example, the rightmost vertex of the triangle 14 forms a separating plane with the leftmost tile edge. This plane intersects the lens area, and divides it into two half-spaces: the negative half-space 16, whose points on the lens cannot see the triangle through any point in the tile, and the positive half-space 18, whose points on the lens potentially can see the triangle. The normal of the separating plane is on the same side of the plane as the tile itself, and this defines the positive and negative half-space on the lens. On the right in FIG. 2, a separating plane is formed from a triangle edge and a tile corner.

Note that we choose the "sign" of the normal (i.e., $+n$ or $-n$) of the separating plane differently depending on whether the triangle vertex, forming the separating plane, is in front or behind the focus plane. The rationale for this is that we would like to cull regions where $h_i(u, v) < 0$, independent of vertex position. For vertices behind the focus plane, the separating plane's normal is chosen such that its positive half space contains the entire tile. In contrast, for vertices in front of the focus plane, the separating plane's normal is such that its negative half-space contains the entire tile. This is illustrated in FIG. 1.

The direction of the two-dimensional normal of the half-space line is inherited from the corresponding separating plane's normal. These two-dimensional normals are illustrated as arrows in FIG. 2. Geometrically, it is easy to see that no point in the negative half-space on the lens can "see" the triangle through any point in the tile.

3

The second set of half-space lines are generated from separating planes formed from a triangle edge and a tile corner. An example is illustrated to the right in FIG. 2. We denote these by Π_j to distinguish them from the planes (π_i) formed by tile edges and triangle vertices. The Π_j planes also generate half-space lines, which are denoted $H_j(u,v)=0$.

The first set of half-space lines, h_i , will be either horizontal or vertical, and in general, the second set of half-space lines, H_j , can have arbitrary orientation. When all tile edges each generate one separating plane, they will form a two-dimensional box in the plane of the lens. With such a box, it is simple and efficient to cull large portions of the samples in the entire tile from further processing. An example is shown to the left in FIG. 3. To the right in the same figure, further regions on the lens have been culled away by the H_j planes. When all triangle vertices either are in front of or behind the focus plane, the remaining active region on the lens is defined by a convex region, where $h_i(u,v) \geq 0$ and $H_j(u,v) \geq 0$. Only the samples with their lens coordinates, (u,v) , inside this active region (region 20 in FIG. 3), need to be inside-tested. The efficiency of our practical algorithms, described later, stems from this fact.

In FIG. 3, the lens is the thick-lined square. On the left, the first four separating planes generate horizontal and vertical half-space lines, which are defined by $h_i(u,v)=0$. Together, they form a two-dimensional bounding box in the lens plane. Only the samples with lens positions in the region 20 need to be further processed. Using only the first four planes can cull away a substantial region of the lens. On the right in FIG. 3, the second set of half-space lines, $H_j(u,v)=0$ can further reduce the region 20. When the triangle is crossing the focus plane, the situation is more complex and we will show hereinafter how to use the first set of half-space lines, $h_i(u,v)$, for culling in such cases.

It is not always possible to create separating planes. For example, when the triangle cuts through a tile, it will not be possible to find separating planes for that tile. However, this occurs when the triangle goes from front- to backfacing or vice versa, and hence, it is possible to add a half-space line to cut away parts of the lens where the triangle is backfacing. This could be beneficial for triangles close to the silhouette of the object. In practice, there is usually a small number of such triangles compared to the rest of the triangles in a scene, and as a consequence, we have not found that this pays off.

We will use an abundant notation in order to simplify the description of the algorithm. As before, the triangle vertices are called p^0, p^1, p^2 . The four corners of a tile are denoted by t^a, t^b, t^c and t^d . The projection of a vertex, p^j , through a tile corner, t^i , onto the lens will be denoted by $l^{ij}=(l_u^{ij}, l_v^{ij})$.

Referring to FIG. 4, showing an illustration of our notation, the vertices, p^0, p^1 , and p^2 , of a triangle are projected through the tile corners, t^a, t^b, t^c , and t^d , onto the lens to form the lens coordinates, l^{ij} , where $i \in \{a,b,c,d\}$, and $j \in \{0,1,2\}$. For example, the vertex p^2 is projected through t^b to generate l^{b2} .

The projection of a vertex, p^j , through a tile corner, gives the lens coordinates of l^{ij} :

$$l_u^{ij} = -\frac{p_x^j F}{p_w^j - F} + t_x^i \frac{p_w^j}{p_w^j - F} = o_u^j + t_x^i \delta^j, \quad (1)$$

$$l_v^{ij} = -\frac{p_y^j F}{p_w^j - F} + t_y^i \frac{p_w^j}{p_w^j - F} = o_v^j + t_y^i \delta^j.$$

The offsets, o , and the deltas, δ , are constant when rendering a particular triangle, and these can therefore be computed in a triangle setup. In addition, there is only a linear dependence

4

on the tile coordinates, (t_x^i, t_y^i) , implying that they can be evaluated efficiently using incremental updates when moving from one tile to the next. If $p_w^j - F = 0$, then that coordinate will never generate a half-space on the lens that will actually cull anything. The reason is that the projections, (l_u^{ij}, l_v^{ij}) , will approach $\pm\infty$, which will never help in culling on a finite lens. Hence, such vertices may be ignored for the remainder of the computations.

Next, we describe how the separating planes, π , through a tile edge and a triangle vertex are found. The most straightforward way to determine whether a plane is separating, is to test whether the two triangle vertices, which were not used to define the plane, are on the opposite side of the plane compared to the tile. We will describe one efficient way of doing this.

Since these half-space lines, $h(u,v)=0$, will be either horizontal or vertical, all computations can be done in two dimensions. Let us assume that we want to determine whether a plane from one of the triangle vertices, p^j , is a separating plane through a tile edge located at $x=t_x^i$. In this case, all calculations can be done entirely in the xw -plane. In the following, recall that the focus plane is at $w=F$, and the lens is at $w=0$. Note also that in all our illustrations, the pixels are located exactly in the focus plane.

To determine whether $u=l_u^{ij}$ actually defines a half-space line from a separating plane, we want to determine in which half-space the other two triangle vertices are located. If they both are not in the same half-space as the tile itself, we have found a separating line, $u=l_u^{ij}$.

To test if $u=l_u^{ij}$ is a separating plane, we set $q=p^j$ and let r be one of the other triangle vertices. We derive a two-dimensional line equation from (q_x, q_w) to (t_x^i, F) , and insert the two other triangle vertices into the line equation. After a bit of mathematical manipulation, we find that the line equation evaluation at a point (r_x, r_w) is:

$$e(q, r) = r_x(q_w - F) + q_x(F - r_w) + t_x^i(r_w - q_w) = O_{qr} + t_x^i \Delta_{qr}. \quad (2)$$

Note also that $e(q,r) = -e(r,q)$, so for a given tile, only $e(p^0, p^1)$, $e(p^1, p^2)$, and $e(p^2, p^0)$ need to be evaluated. In general, for $u=l_u^{ij}$ to define a separating line, the two other triangle vertices should be on the same side of the line, and be in the negative half-space when $p_w^j > F$, and in the positive half-space when $p_w^j < F$, as can be seen in FIG. 1. The case when $p_w^j - F = 0$ is ignored, as described previously, because such projections will not provide any culling. Note that Equation 2 is linear in t_x^i , which is to be expected since l_u^{ij} is linear.

For example, given the vertex $q=p^0$, and the left-most tile edge, $x=t_x^a$, we test if $u=l_u^{a0}$ is a separating half space-line by evaluating the line equation (Equation 2) for $r=p^1$ and $r=p^2$. If it is separating, the corresponding half-space line, $h(u,v)$, is defined by:

$$h(u, v) = \begin{cases} u - l_u^{a0}, & \text{when } p_w^0 > F, \\ l_u^{a0} - u, & \text{when } p_w^0 < F, \end{cases} \quad (3)$$

which is a vertical half-space line on the lens. Note that the culling "direction" changes depending on whether the vertex is in front of the focus plane, F , or behind it. In addition, the $p_w > F$ tests are reversed when testing against the right-most tile edge, $x=t_x^d$. Similar equations are created for all lens coordinates, l_u^{ij} and l_v^{ij} that have been verified to be real separating planes (using Equation 2). This is all that is needed

5

for computing the horizontal and vertical half-space lines, $h_i(u,v)=0$.

Finally, we describe how we handle cases where at least one vertex is located in front of the focus plane, and at least one behind the focus plane.

In FIG. 6, when a triangle intersects focus plane, a tile side can generate two separating planes. Note that culling is done in the “inner” region 24 in this case, in contrast to when the triangle does not intersect the focus plane (see FIG. 1). In this case, we generate one uv-box for the topmost region 25 and one for the bottommost region on the lens.

The set of triangle vertices behind focus plane is denoted A, and the set in front of the focus plane is denoted B. When one set is empty, all vertices are located on one side of the focus plane. In these cases, it is straightforward to generate a box in the uv domain on the lens, and inside-testing will only be done inside the box. We call these uv-boxes, where we use the term box broadly because they can extend infinitely in some directions as we will see. When both A and B are non-empty, and there are separating half-space lines generated from both sets, two uv-boxes will be generated. This can be seen in FIG. 6. When there are two separating half-space lines generated in one set (e.g., A), both these will be used to define that set’s uv-box.

This is illustrated in FIG. 5. Visible regions for the triangle seen through the tile 23. In this case, we generate two uv-boxes on the lens, and inside testing is only done for samples with uv-coordinates inside these regions.

After computing l^j , the classification of planes can be performed in a different way that requires less live data and is potentially more amenable to parallelization. The algorithm below is only described for the u dimension; the same procedure is valid for the v dimension.

We start by classifying the triangle vertices depending on which side of the focus plane they are. Using the notation from FIG. 4, we then form the two intervals:

$$\hat{A} = [A, \bar{A}] = \left[\min_{i,j} l_{i,j}^j, \max_{i,j} l_{i,j}^j \right] \forall i, j: i \in \{a, d\}, p_w^j > F,$$

$$\hat{B} = [B, \bar{B}] = \left[\min_{i,j} l_{i,j}^j, \max_{i,j} l_{i,j}^j \right] \forall i, j: i \in \{a, d\}, p_w^j < F. \quad (4)$$

If $p_w^j > 0$, Equation 4 can be simplified, and we only need to test one of the two tile corners for each interval limit. If all triangle vertices are on one side of the focus plane, one of the intervals is empty, and the visible interval in u is given by \hat{A} or \hat{B} . If $\hat{A} \cap \hat{B} \neq \emptyset$, no separating plane exists between a vertical tile edge and a triangle corner.

Finally, if $\hat{A} \neq \emptyset$, $\hat{B} \neq \emptyset$ and $\hat{A} \cap \hat{B} \neq \emptyset$, we get the visible intervals:

$$[-\infty, \bar{A}], [\bar{B}, \infty] \text{ if } \hat{A} < \hat{B},$$

$$[-\infty, \bar{B}], [\bar{A}, \infty] \text{ if } \hat{B} < \hat{A}. \quad (5)$$

The intervals for the v dimension are computed similarly. To derive the visible 2D regions on the lens, the intervals including A borders for u and v are combined to form a 2D region. Similarly for the intervals including B borders. For example, if we get the case described in Equation 5 in both u and v, with $\hat{A}_u > \hat{B}_u$ and $\hat{A}_v > \hat{B}_v$, the resulting uv-boxes are:

6

$$uv\text{-box } A: [\bar{A}_u, \infty] \times [\bar{A}_v, \infty] \quad (6)$$

$$uv\text{-box } B: [-\infty, \bar{B}_u] \times [-\infty, \bar{B}_v]. \quad (7)$$

This case is illustrated in FIG. 5.

The algorithm described so far is correct as long as no vertex lies exactly in the focus plane. A few minor adjustments need to be performed on the calculated regions whenever such vertices exist:

$$\bar{A}_u = \infty, \bar{B}_u = -\infty, \text{ if } p_x^j \leq t_x^d \text{ for any } p^j \text{ with } p_w^j = 0,$$

$$\bar{A}_u = -\infty, \bar{B}_u = \infty, \text{ if } p_x^j \geq t_x^a \text{ for any } p^j \text{ with } p_w^j = 0, \quad (8)$$

and similarly for v.

Next, we describe how the H_j lines are computed.

Given a triangle edge qr, we form the ray $q+t(r-q)=q+td$. Referring to FIG. 7, we compute the intersection point, v, between the ray and the focus plane. We divide the focus plane into nine regions and identify which region v falls into. The intersection point, v, does not depend on the tile for which we wish to categorize the edge. It is therefore possible to pre-compute this point and at which tile the cell changes. The only work needed per tile is thus a comparison of the tile coordinate with these pre-computed boundaries.

With v categorized into one of the nine cells, we identify the two tile corners m and n which form the largest angle $\angle t^m v t^n$. These can be tabulated as follows:

Cell	1	2	3	4	5	6	7	8	9
m	b	b	d	a	—	d	a	c	c
n	c	a	a	c	—	b	d	d	b

Using this table, we can form two candidate planes $\Pi_1: (q, r, t^m)$ and $\Pi_2: (r, q, t^n)$. Using the sign of the d_w -component, i.e., whether the edge, qr, points towards the camera, we can choose plane normal directions such that the tile is in the negative halfspace of the respective plane. For edges parallel to the image plane, $d_w=0$, and we use d_x, d_y to determine m and n and the sign of $q_w - F$ to determine the normal direction. There are no candidate planes for cell five, since in this cell, the edge is pointing straight at the tile and there cannot exist any separating planes between the edge and a tile corner. Likewise, any edge qr where $q_w = r_w = F$, cannot produce a useful half-space line and is thus ignored.

In FIG. 7, the triangle edge qr is intersection tested against the focus plane, which generates the intersection point, v. The tile abcd divides the focus plane into nine regions, numbered 1-9. In the example above, the intersection point, v, lies in cell 1, and therefore the candidate corners to form separating planes are b and c. In this illustration, those are shown as lines, from v to c and from v to b which form tangents between v and the tile.

To determine if Π_1 and Π_2 are separating planes, the third vertex of the triangle is tested against these planes. If the vertex is in the positive half-space of a plane, that plane is a separating plane. Each triangle edge can generate up to two separating planes, as can be seen in FIG. 8.

In FIG. 8, the two candidate planes, Π_1 and Π_2 , divide space into four regions, shown as I-IV. By construction, the triangle edge qr is included in both these planes. Whether Π_1 and Π_2 are separating depends on which region the third triangle vertex, s, is in. If s is in region I or II, Π_1 is a separating plane. Π_2 is separating if s lies in region II or III. A triangle edge qr can thus produce zero, one, or two separating planes.

Given a plane $\Pi: n \cdot x + c = 0$, the corresponding half space line on the lens (u,v,0) is

$$n \cdot (u, v, 0) + c = n_x u + n_y v + c. \quad (9)$$

7

This equation varies non-linearly and must therefore be computed for each tile. To see this, consider again the triangle edge qr and a tile corner t in the focus plane. A plane through these points is defined by $\mathbb{P}:n \cdot x + c = 0$, where:

$$n = q \times r + (r - q) \times t$$

$$c = -t \cdot (q \times r) \quad (10)$$

Note that the normal n and c change as we move in screen space. The change is proportional to the difference in depth of q and r .

Finally, the normals of the planes Π_i need to be handled carefully when culling samples on the lens. When culling regions from uv-box A , samples in the positive half-space of Π_i can be culled. However, when culling regions from uv-box B , samples from the negative half-space of Π_i can be culled.

Up until now, the half-space lines, h_i , and H_j , which lie in the plane of the lens, have been computed in an exact manner. In this subsection, we will describe how they can be exploited to quickly cull samples for faster depth of field rasterization. The half-space lines are computed on a per-tile basis, and hence culling opportunities will be shared between all the pixels in a tile. However, the process still needs to be very efficient if we are to see performance gains. As will be seen, determining which samples lie within the active subspace of the lens is essentially a rasterization process in itself.

We superimpose a square grid on top of the lens shape and keep track of the number of samples falling into each grid cell. This is illustrated in FIG. 9. These sample distributions vary from pixel to pixel, and we use a small set (32×32) of distributions that are scrambled and repeated over the screen. Note that 32×32 can be replaced with a resolution of the implementor's choice.

FIG. 9 shows a lens grid with corresponding storage. For each grid cell, we store the number of samples within that cell, and an offset pointing to the first sample in that grid cell. With this layout, we can efficiently cull large sets of samples against the separating half-space lines.

The half-space lines, h_i , which are generated from triangle vertices and tile edges, provide an easy means of culling since they are axis aligned. We can simply clamp them down and up to the nearest grid cell and use the clamped rectangular extents to quickly traverse relevant samples.

For the non-axis aligned half-space lines, H_j , we iterate over all grid cells in the rectangular extents computed from the h_i lines and conservatively test for overlap between the grid cells and the H_j lines. This essentially boils down to a micro-rasterization process for every screen-space tile in order to cull the samples. A way to optimize this is to exploit the limited resolution of the lens grid and use a pre-computed rasterizer. The idea here is to have a small lookup-table, where inside bitmasks are precomputed and stored for different orientations of edges (e.g., a 64-bit mask table is sufficient for an 8×8 grid). All edges are queried in the table, and the inside region is computed by doing AND-operations on the bitmasks.

In the following pseudo-code, we outline the full algorithm for rendering depth of field using half-space line culling.

```

compute BBox of defocused triangle Δ
compute initial  $h_0, \dots, h_3$  half-space lines
for all tiles  $T$  in BBox do
  inexpensive update of  $h_0, \dots, h_3$ 
  compute UVBoxes from  $h_0, \dots, h_3$ 
  if any UVBox overlaps lens shape then
    compute  $H_j$  half lines
    bitmask = compute overlap bitmask for all  $H_j$ 

```

8

-continued

```

for all pixels  $p$  in  $T$  do
  for all grid cells  $C$  in UVBoxes do
    if bitmask  $[C] == \text{True}$  then
      test samples in  $C(p)$  against  $\Delta$ 
    end if
  end for
  shade pixel (multisampling)
end for
end if
end for

```

According to one embodiment, a high-level depth of field rasterization algorithm is shown in FIG. 10. The sequence shown in FIG. 10 may be implemented in software, hardware, and/or firmware. In a software embodiment, the sequence may be implemented by computer readable instructions stored on a non-transitory computer readable medium, such as an optical, semiconductor, or magnetic storage device.

Initially, a tile is selected (block 28) and the separating planes are developed between a triangle to be rendered and the tile being processed, as indicated in block 30. Next, half-space regions on the lens are defined using the separating planes, as indicated in block 32. Then, the useful half-space regions are identified, as indicated in block 34. The useful half-space regions are those regions on the lens from which a triangle to be rendered is visible. The non-useful half-space regions are culled, as indicated in block 36, and are not further processed. As indicated in block 38, only the useful half-space regions are inside tested. In block 40, the half-space region lines are incrementally updated. A check at diamond 42 determines whether this is the last tile. If so, the flow ends and, otherwise, the flow iterates back to block 28 to select another tile to process.

The computer system 130, shown in FIG. 11, may include a hard drive 134 and a removable medium 136, coupled by a bus 104 to a chipset core logic 110. The computer system may be any computer system, including a smart mobile device, such as a smart phone, tablet, or a mobile Internet device. A keyboard and mouse 120, or other conventional components, may be coupled to the chipset core logic via bus 108. The core logic may couple to the graphics processor 112, via a bus 105, and the central processor 100 in one embodiment. The graphics processor 112 may also be coupled by a bus 106 to a frame buffer 114. The frame buffer 114 may be coupled by a bus 107 to a display screen 118. In one embodiment, a graphics processor 112 may be a multi-threaded, multi-core parallel processor using single instruction multiple data (SIMD) architecture.

In the case of a software implementation, the pertinent code may be stored in any suitable semiconductor, magnetic, or optical memory, including the main memory 132 (as indicated at 139) or any available memory within the graphics processor. Thus, in one embodiment, the code to perform the sequence of FIG. 10 may be stored in a non-transitory machine or computer readable medium, such as the memory 132, and/or the graphics processor 112, and/or the central processor 100 and may be executed by the processor 100 and/or the graphics processor 112 in one embodiment.

FIG. 10 is a flow chart. In some embodiments, the sequences depicted in these flow charts may be implemented in hardware, software, or firmware. In a software embodiment, a non-transitory computer readable medium, such as a semiconductor memory, a magnetic memory, or an optical memory may be used to store instructions and may be executed by a processor to implement the sequences shown in FIG. 10.

The graphics processing techniques described herein may be implemented in various hardware architectures. For example, graphics functionality may be integrated within a chipset. Alternatively, a discrete graphics processor may be used. As still another embodiment, the graphics functions may be implemented by a general purpose processor, including a multicore processor.

References throughout this specification to “one embodiment” or “an embodiment” mean that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one implementation encompassed within the present invention. Thus, appearances of the phrase “one embodiment” or “in an embodiment” are not necessarily referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be instituted in other suitable forms other than the particular embodiment illustrated and all such forms may be encompassed within the claims of the present application.

While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

What is claimed is:

1. A method of depth of field rasterization comprising:
parameterizing a lens area, having regions, based on its shape;
identifying lens regions in which a convex polygon to be rendered is not visible; and
culling, using a graphics processor, regions on a lens from which a convex polygon to be rendered is not visible.
2. The method of claim 1 including only performing inside tests on the remaining, uncullied regions.
3. The method of claim 2 including culling half-space regions on a lens from which a triangle is not visible.
4. The method of claim 1 including using separating planes between a convex polygon to be rendered and a tile being processed to partition the lens into regions, where a subset of the regions is inside tested.
5. The method of claim 4 including using separating planes and for each separating plane, defining a half-space line on the lens plane.
6. The method of claim 5 including processing only samples inside the remaining uncullied regions.
7. The method of claim 4 including developing said half-space regions from separating planes formed from a triangle vertex and a tile edge.
8. The method of claim 4 including developing said half-space regions from separating planes formed from a triangle edge and a tile corner.
9. The method of claim 8 including locating said separating planes by intersection testing a triangle edge against a focus plane to generate an intersection point, dividing the focus plane into regions and identifying regions that the intersection point falls within.
10. The method of claim 7 including incrementally updating half-space lines between half-space regions.
11. A non-transitory computer readable medium storing instructions executed by a computer to:
parameterize a lens area, having regions, based on its shape;
identify lens regions in which a convex polygon to be rendered is not visible; and

cull regions on a lens from which a convex polygon to be rendered is not visible for depth of field rasterization.

12. The medium of claim 11 further storing instructions to only perform inside tests on the remaining, uncullied regions.

13. The medium of claim 12 further storing instructions to cull half-space regions on a lens from which a triangle is not visible.

14. The medium of claim 11 further storing instructions to use separating planes between a convex polygon to be rendered and a tile being processed to partition the lens into regions, where a subset of the regions is inside tested.

15. The medium of claim 14 further storing instructions to use separating planes and for each separating plane, define a half-space line on the lens plane.

16. The medium of claim 15 further storing instructions to process only samples inside the remaining uncullied regions.

17. The medium of claim 14 further storing instructions to develop said half-space regions from separating planes formed from a triangle vertex and a tile edge.

18. The medium of claim 14 further storing instructions to develop said half-space regions from separating planes formed from a triangle edge and a tile corner.

19. The medium of claim 18 further storing instructions to locate said separating planes by intersection testing a triangle edge against a focus plane to generate an intersection point, dividing the focus plane into regions, and identifying regions that the intersection point falls within.

20. The medium of claim 17 further storing instructions to incrementally update half-space lines between half-space regions.

21. An apparatus comprising:
a storage; and
a processor coupled to the storage to parameterize a lens area, having regions, based on its shape, identify lens regions in which a convex polygon to be rendered is not visible, and perform depth of field rasterization by culling regions on a lens from which a convex polygon to be rendered is not visible.

22. The apparatus of claim 21, said processor to only perform inside tests on the remaining, uncullied regions.

23. The apparatus of claim 22, said processor to cull half-space regions on a lens from which a triangle is not visible.

24. The apparatus of claim 21, said processor to use separating planes between a convex polygon to be rendered and a tile being processed to partition the lens into regions, where a said set of the regions is inside tested.

25. The apparatus of claim 24, said processor to use separating planes and for each separating plane, to define a half-space line on the lens plane.

26. The apparatus of claim 25, said processor to only process samples inside the remaining uncullied regions.

27. The apparatus of claim 24, said processor to develop said half-space regions from separating planes formed from a triangle vertex and a tile edge.

28. The apparatus of claim 24, said processor to develop said half-space regions from separating planes formed from a triangle edge and a tile corner.

29. The apparatus of claim 28, said processor to locate said separating planes by intersection testing a triangle edge against a focus plane to generate an intersection point, divide the focus plane into regions, and identify regions that the intersection point falls within.

30. The apparatus of claim 27, said processor to incrementally update half-space lines between half-space regions.